

### The Kinetica Advantage

#### Performance

- Ingest streaming data—billions of records per minute—and get “up to the moment” analytics
- Realize 100x performance improvement on queries compared to CPU-based in-memory solutions
- Holds 100s of TB of data in-memory for extremely low-latency analytics

#### Simplicity

- No typical tuning or indexing required; ask and answer any question in real time
- Connect with common BI tools like Tableau, Kibana and Caravel
- A converged, unified suite; not multiple disparate components

#### Predictably Scalable

- Easily scale up or out
- Data written to Kinetica is automatically routed to parallel connections across the cluster
- OLAP queries are executed using fully distributed GPU-accelerated processing across the cluster

#### Easy APIs and Integration

- Open source integration components include Apache NiFi, Spark and Spark Streaming, Storm, Kafka and Hadoop
- Kinetica's APIs are fully supported in REST, Java, Python, C++, Javascript and Node.js
- ODBC and JDBC drivers integrate with industry-standard BI and SQL tools

#### Complete Native Visualization and Geospatial Capabilities

- Real-time geoprocessing
- A fully GPU-accelerated distributed rendering pipeline
- Visualize billions of points in seconds

## IBM Power Architecture with NVLink provides the most advanced platform for Kinetica's high-performance analytics.

IBM's Power architecture with the new NVLink interface provides the most advanced and most affordable platform for high performance analytics with Kinetica today. The new NVLink interconnects seen in IBM's S822LC 'Minsky' machines open a much wider path between the CPU and GPU, enabling Kinetica to take full advantage of system memory.

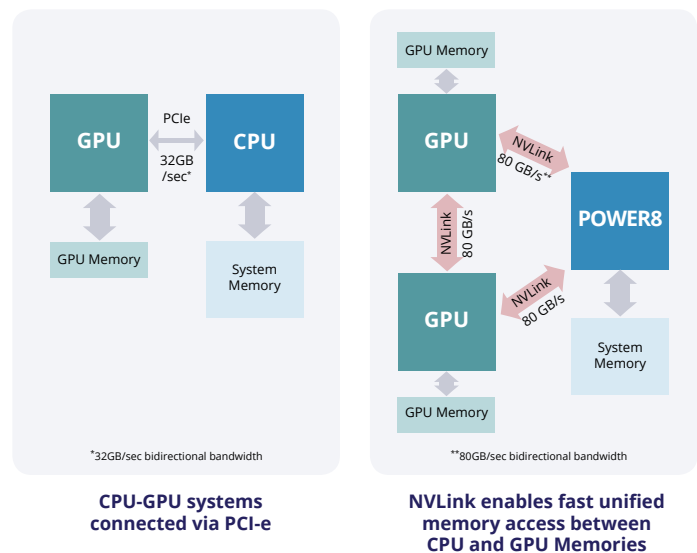
Data analytics applications powered by Kinetica, have often been limited by the bandwidth of the PCI Express interface that connects CPUs and GPUs in today's servers. For Kinetica, which is designed to take full advantage of system memory, swift movement of data between CPU and GPU is absolutely critical to getting the best performance for analytics on large datasets.

### Power Architecture and NVLink

Previously, a typical high performance compute system today would have one or more GPUs connected to a CPU using PCI Express. Even at the fastest PCIe 3.0 speeds, the bandwidth provided over this link pales in comparison to the bandwidth available between the CPU and its system memory. In a multi-GPU system, the problem is compounded with the limited PCIe bandwidth being shared between the GPUs. The resource contention gets even worse when peer-to-peer GPU traffic is factored in.

### 2.5x Faster CPU-GPU Connection via NVLink

The NVLink interface opens the full potential of GPU accelerators and large scale analytics with Kinetica, by enabling brute-force processing to be done on both the GPU, while the CPU handles tasks it's well designed for. For example, simple database operations can be done on the CPU, but “sorting” algorithms can be offloaded to the GPU accelerator by quickly transferring the data to and fro from the CPU to the GPU.



### Real-time Analytics with Kinetica's GPU-Accelerated Database

Out think the status quo. Be able to better understand your customers and your business by querying massive datasets in seconds vs. hours to help you become a real-time retail enterprise.

---

## Solve Precision Marketing with Kinetica, IBM and NVIDIA

Meet demand spikes in real time, and convert billions of impressions to millions of dollars. Out-think out of stock with Kinetica on IBM's OpenPOWER LC servers—the accelerated platform of choice leveraging NVIDIA NVLink technology. This is your solution to precision marketing.

---

### Recommended Hardware



**S822LC “Minsky” for  
High Performance Computing**

#### GOLD CONFIGURATION

CPU Cores	2 x 10 core POWER8 with NVLink CPUS, 2.86Ghz
RAM	1 TB Memory
SSDs/Flash Disk	2 x 1TB SSD (for OS and local storage), 3.2 TB NVMe
GPUs	Bidirectional, point-to-point between GPU-CPU and GPU-GPU
GPU Cores	3584 per GPU
GPU VRAM	16 GB per GPU

#### SILVER CONFIGURATION

CPU Cores	2 x 8 core POWER8 with NVLink CPUS, 3.2Ghz
RAM	1 TB Memory
SSDs/Flash Disk	2 x 1TB SSD (for OS and local storage), 3.2 TB NVMe
GPUs	Point-to-point between GPU-CPU and GPU-GPU
GPU Cores	3584 per GPU
GPU VRAM	16 GB per GPU



**For more information on Kinetica and GPU-accelerated databases, visit [kinetica.com](http://kinetica.com)**

Kinetica and the Kinetica logo are trademarks of Kinetica and its subsidiaries in the United States and other countries. Other marks and brands may be claimed as the property of others. The product plans, specifications, and descriptions herein are provided for information only and subject to change without notice, and are provided without warranty of any kind, express or implied. Copyright © 2016 Kinetica